Formalisation of Safety Reasoning in Protocols and Hazard Regulations

Peter Hammond, Brunel University, Uxbridge, UK
Marek J Sergot, Imperial College of Science, Technology and Medicine, London, UK
Jeremy C Wyatt, Imperial Cancer Research Fund, London, UK.

Written protocols are often employed to guide patient care. For treatment within a clinical trial, compliance with the trial protocol may be critical in ensuring efficacy and safety. Previous empirical work has established generic safety principles for reasoning about adverse events in clinical trials and their formalisation has been applied in a decision support system for managing treatment plans in oncology. The same generic knowledge can be reused to generate specific safety clauses when designing new treatment plans. Typically, clinicians devise trial protocols relatively infrequently and so software aids, especially those assisting with regulatory/safety conformance, will encourage more effective use of their time. A similar approach to the formalisation of safety knowledge in the control of hazardous industrial processes is discussed.

INTRODUCTION

Protocol-based care is increasingly popular in medicine. Primarily, protocols and guidelines are employed to standardise some aspect of treatment: to regularise the management of a disease or to enforce a particular protocol during a clinical trial of a new therapeutic agent. Some guidelines do not need to be followed rigorously, whereas adherence to a clinical trial protocol is essential if statistical analysis of the trial results is to be scientifically valid or if the safety of patients is to be ensured. Compliance is also an important issue for pharmaceutical companies when a clinical trial involves the evaluation of a new product in anticipation of its approval by bodies like the US Food and Drug Agency before its commercialisation.

The use of computers in the application of clinical guidelines and protocols is becoming more widespread and improved compliance[1] and more complete data capture for subsequent analysis of clinical trial results[2] have been reported. In particular, decision support systems for protocol-based care in oncology have been studied for more than 20 years. The safety aspects of such systems were studied in an earlier project and some of the resulting generic "safety principles"[3] were encoded in OaSiS[4], a prototype knowledge-based system supporting protocol-based care in cancer management. A small number of computer aids for designing new clinical trials, therapy plans and associated protocol documents have been implemented. Notable examples are DaT[5] and OPAL[6].

We summarise previous work on the empirical derivation, formalisation in first-order logic and partial implementation of safety reasoning for cancer management. Then we describe the reuse of this safety knowledge to generate specific safety clauses semi-automatically during clinical trial design and protocol document assembly. Finally, the animation of codes of practice for regulating the use of hazardous substances in industrial processes is shown to be a promising new domain for similar analysis and formalisation.

METHODS & RESULTS

General Safety Principles in Oncology

In clinical trials of cancer treatments, combinations of surgery, hormone-, radio- and chemotherapy are evaluated for their efficacy, toxicity and cost. The complexity of associated therapy plans and protocols and the volume of data collection demands computer support. Besides treatment recommendations, protocols also regulate modifications to treatment resulting from toxic side-effects.

Following a detailed survey of more than 50 cancer trial protocols and discussions with clinicians, particular instances of safety reasoning about adverse events were reformulated as a collection of generic principles. These general safety principles are in addition to the obvious curbs on drug administration in terms of maximum single, course and lifetime cumulative dosages. Each of the nine principles is given in figure 2 along with an illustrative example.

The summary form of the principles and the rule-based form illustrated below in figure 1 were always stated in a generic format in anticipation of their use in other application domains. These were either culled from the oncology literature or identified in discussions with clinical oncologists and pharmacists.

As Figure 2 demonstrates, many of the examples found required detailed discussion with oncologists and pharmacists in order to uncover the implicit clinical and pharmacological reasoning involved. Once the generic principles were identified, they were expressed informally as rules and subsequently formally in first order logic[3]. For example, the exacerbation principle can be represented informally in the form of the following rule

Action1 should not be performed during Action2 in Plan if
  Action2 is necessary part of Plan and
  Action2 produces or may produce Effect and
  Effect is potentially hazardous and
  Action1 aggravates or makes Effect more likely and
  Action1 has alternative in Plan not aggravating Effect

Figure 1: Informal Version of Exacerbation Principle

Warning	Warn about hazards due to inadequate execution of essential actions
	Slow infusion of piroxantrone is required because of risk of major motor seizures[7]
Reaction	React appropriately to ameliorate detected hazards
	Administer diuretics and digoxin and withdraw doxorubicin if signs of cardiac failure are noticed[8]
Exacerbation	Avoid exacerbating anticipated hazards
	Nephrotoxic antibiotics such as Gentamicin should be avoided during and immediately after the Cisplatin infusion[8]
Monitoring	Monitor responses which herald hazardous situations
	Measurement of methotrexate levels is essential ... at 24 and 48 hours (after its administration)[8]
Efficacy	Ensure that overall plans are efficacious in pursuing stated objectives
	Tumour necrosis factor is not a useful antitumour agent if administered IV or IM[9]
Sequencing	Order (essential) actions temporally for good effect and least harm
	Taxol is given before cisplatin. The reverse order can produce severe neutropenia[10]
Diminution	Avoid undermining the benefits of essential actions
	Aspirin reduces efficacy of Interferon-a2α.[11]
Critiquing	Critique the proposal of certain hazardous actions even if they are well motivated
	Doses of etoposide should not be reduced for elevated serum bilirubin concentrations[7]
Prevention	Prevent or ameliorate hazards before executing an essential action
	Folinic acid rescue helps ameliorate methotrexate-induced bone marrow suppression[8]

Figure 2: Safety principles and oncology examples

In the OaSiS system, we use a logic database model[12] to represent the treatment plan negotiated from the system's recommendations and the user's amendments. The prohibitions in the safety rules (for example, those corresponding to the exacerbation and diminution principles) act as integrity constraints on updates to this treatment database. Thus, any attempt to amend system generated therapy suggestions or to introduce new therapies can be challenged if it contravenes these prohibitions. A formalisation of the exacerbation principle is given in figure 3.

Although it is usual to write integrity constraints in the form of denials, we find it more convenient to employ the form shown. The reason is simply that all conditions except for *user_suggestion* are static, in the sense that they are stored in parts of the OaSiS system which the user does not modify during a consultation. Here *invalid* can be read as an alternative symbol to ¬ for (standard, truth-functional) negation. It also has an operational meaning, to signal that in the case of violation of the constraint it is this condition, i.e., the attempted input, which is to be rejected.

Label	Clause
rule_a	invalid(user_suggestion(perform(Action1),Plan))← part_of(Action2, Plan) ∧ produces_effect(Action2, Effect) ∧ hazardous(Effect) ∧ aggravates(Action1, Effect) ∧ is_avoidable(Action1, Effect, Plan).

Figure 3: Formal Version of Exacerbation Principle

To reproduce the exacerbation example, we need to add specific clauses as defined in figure 4

Label	Clause
e	produces_effect(admin_of(cisplatin),uraemia).
h	hazardous(uraemia).
a	aggravates(admin_of(Drug),uraemia) ← drug_type (Drug, nephrotoxic_anti_biotic).
d	drug_type(gentamicin,nephrotoxic_anti_biotic).

Figure 4: Formalisation of Exacerbation Example

Any attempt to prescribe gentamicin will contravene the constraint in figure 4 and can be challenged.

The underlying logical model is straightforward to describe. Suppose the knowledge base has the following components

Oncology	general oncology knowledge
Safety	the generic safety principles
Protocol	a particular protocol
Treatment	individual patient treatment plan negotiated between system and clinical user.

Notice that *Oncology*, *Safety* and *Protocol* are static whereas *Treatment* alters as the patient is treated over the course of the clinical trial. We have {rule_a, h} ⊆ Safety and {e, a, d} ⊆ Oncology. Additional knowledge about the patient's condition, *Observations*, is presented on the fly during a consultation with OaSiS. We should record any *Suggestions* from the clinical user which attempt to amend the protocol-derived therapy or which involve new treatments, for example for conditions other than the malignancy. Let Clinician = Observations ∪ Suggestions. Then, the treatment plan is deduced from the knowledge base components *Oncology*, *Protocol* and *Clinician* subject to the validity of the integrity constraints expressed in *Safety*, i.e.,

(1) Oncology ∪ Protocol ∪ Clinician ⊢ Treatment

subject to the consistency of

Oncology ∪ Protocol ∪ Clinician ∪ Safety.

Drafting safety clauses in oncology protocols

Consider now the situation where a clinician or a pharmaceutical company wishes to design a new clinical trial and draft the associated protocol document. In particular, the protocol needs to include specific safety clauses covering the treatments to be administered in the trial. We no longer need any patient-specific information, namely *Clinician* and *Treatment* , but we do we need to introduce:

TherapyPlan	the protocol designer's description of the intended therapy regimes (drugs, dosages, routes of administration, cycles of therapy etc.)

The idea is to use *Oncology* and *Safety* along with *TherapyPlan* to deduce symbolic representations of parts of the paper protocol and at the same time generate software components that can be used with OaSiS to help manage the therapy regime described in the new protocol. In particular, the derivation of the particular safety clauses to be included in a new protocol can be expressed as in (1) above:

(2) Oncology ∪ Safety ∪ TherapyPlan ⊢ ProtocolSafety

where we treat the generic safety regulations as part of the overall knowledge base now and use them deductively rather than as a separate set of integrity constraints. Thus, the specific safety clauses for the protocol, *ProtocolSafety*, are the dynamic, deducible component of the knowledge base.

We illustrate part of this protocol design process with an example. If we were devising a trial, bO03 say, which used the cytotoxic drug *cisplatin* we would want to be able to derive safety clauses such as the exacerbation example, i.e. *Nephrotoxic antibiotics such as Gentamicin should be avoided during and immediately after the cisplatin infusion.* The administration of *cisplatin* in a particular protocol can be interpreted as the dynamic addition of a fact supplied by the protocol designer to the database specifying the detailed structure of the therapy i.e., *TherapyPlan*. Once this fact is added to the knowledge base we are able to derive a theorem as a specialisation of the exacerbation rule :

Oncology ∪ Safety ∪ { part_of(admin_of(cisplatin),bO03)} ⊢
∀ Drug { invalid(user_suggestion(admin_of(Drug), bO03))
← drug_type(Drug, nephrotoxic_anti_biotic) ∧
is_avoidable(admin_of(Drug), bO03)}.

The theorem can be paraphrased as "it is invalid for the user to prescribe the administration of a drug if it is an avoidable nephrotoxic drug".

This reuse of safety knowledge can be modelled in PROLOG by a meta-interpreter which unfolds the generic safety rules suitably to produce more specific safety clauses for inclusion in the trial protocol. In order to determine the most appropriate level at which to stop unfolding elements of clauses, the protocol designer's assistance is required. It would, of course, be necessary to provide suitable natural language translations of the unfolded clauses to the designer so that the appropriate level of generality can be identified. For example, in this case we can derive the theorem as stated or we can unfold drug_type(Drug, nephrotoxic_anti_biotic) further to uncover the individual safety examples by instantiating particular nephrotoxic drugs, including gentamicin.

Animating regulations about industrial hazards

The UK Health & Safety Commission publish approved codes of practice[13] governing hazards arising in various industrial processes. A suitable representation of these codes of practice could be used to animate the regulations so that the relevance of the regulations to individual companies could be determined. We illustrate where the generic safety principles may need to be extended before they can be used in this way.

Some of the nine generic safety principles in figure 2 are almost immediately applicable to this new domain. For example, the monitoring of hazards has an obvious overlap and corresponding to the monitoring example we have

(a) When using vinyl_chloride monomer continuous monitoring should be carried out

(b) Carbon disulphide is produced as a vapour in india rubber processing - annual health checks are needed

There are interesting similarities between administering chemotherapy and handling toxic fumigants. For example, Figure 5 indicates how in the oncology domain, the process (chemotherapy) is used to treat a problem (tumour) in some host (a patient). During chemotherapy, for example, we wish to minimise the negative effects on the host (patient) and maximise the desired effects on the problem (tumour). There is an unestablished or negligible effect on the person, or operator, managing the operation (therapy).

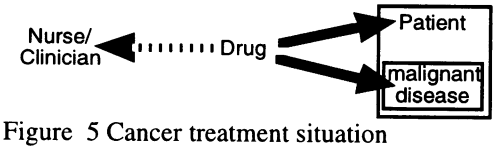

Figure 5 Cancer treatment situation

This is in contrast to some industrial processes (figure 6) where the operation (treatment with a toxic fumigant, say) is used to treat a problem (infestation or contamination) in some host (an industrial building).

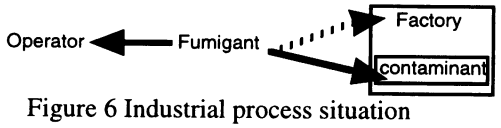

Figure 6 Industrial process situation

The COSHH regulations have little to say about the undesirable effect on the host (the building) but considerably more about the hazards to the industrial

operator or others likely to be in the vicinity during the treatment process. Thus by changing domains, we have quite a different emphasis on the hazards to the individual participants. Indeed, further scrutiny of the COSHH regulations identifies more agents to be taken into consideration, including Employers, Employees, Operators, and members of their families. In addition, besides the more obvious process-related safety regulations, there are detailed descriptions of duties obligating the behaviour of the different participants as individuals and as a group. Thus, in this respect at least, the new domain requires a richer ontology than that arising from the oncology domain analysis.

Finally, an intriguing and interesting similarity is the prohibition of actions which might actually be effective for their intended use but which are disallowed because of their hazardous nature. For example, in the oncology domain, we have a general rule corresponding to the critiquing principle in fig 1:

Plan cannot include Action if
Action ameliorates identified hazardous effect Effect and
Justification is user's motivation for Action and
Justification is not acceptable for ameliorating Effect

Examples of such critiques of motivation are found in a number of cancer protocols, one was given in figure 2 and another is

> No dose modification of r-metHuG-G-CSF based on blood counts will be made[14]

Here, the emphasis is on amendments to the original plan as recommended by the protocol. In the hazardous substances situation we find a similar generic rule

Plan cannot include Action if
Action is effective for Process and
Process is user's motivation for Action and
Action is not acceptable for use in Process

with illustrative examples as follows:

> Sand/free silica cannot be used as an abrasive for blasting
>
> Carbon disulphide cannot be used in the cold-cure process of vulcanising for proofing
>
> Ground flint cannot be used in pottery manufacture in a wash for saggars, trucks, bats or cranks ...

The emphasis in this case is on the activities originally planned by the operator of the industrial process.

Related work

Any clinical computing system which needs to reason about drug therapy is likely to include data on drug interactions, adverse events and the like. The safety or toxicity of drug treatments was a focus in the ONYX project[15] and in the work of Swartout[16]. The ONYX system contains rules which describe drug dose modifications; an illustrative example that is very much in keeping with our approach is

If Drug is a drug in Chemotherapy and Problem is one of the current problems and Drug is not one of causes of Problem and Drug can contribute to Problem
then reduce the dosage of Drug.

Swartout developed the OWL Digitalis Adviser[17] in the late 1970s and honed its explanatory component to considerable sophistication in XPLAIN[18] which used declarative domain models to generate explanations of its problem-solving behaviour. The capabilities of XPLAIN are demonstrated in terms of the toxicity of digitalis therapy:

> The system is anticipating digitalis toxicity. Increased serum calcium causes increased automaticity, which may cause a change to ventricular fibrillation. Increased digitalis also causes increased automaticity. Thus, if the system observes increased serum calcium, it reduces the dose of digitalis due to increased serum calcium.

The work of Musen et al at Stanford University on Protege[19] addresses issues of reuse but in a different manner. Similarly, the critiquing rules in van der Lei's HyperCritic[20] system have much the same flavour as those described here and are potentially reusable.

Our approach differs in its application of logic-based models in the formalisation and implementation of the results in decision support software. We have also tried to record our findings in a much more generic framework so that they are more amenable to application in other domains.

Conclusions and future work

The generic safety knowledge originally determined for the run-time application of protocol-based care in cancer appears to be reusable for the process of designing new clinical trials and protocols governing treatment. There is still work to be done to prove the approach can be successful in a full implementation of a knowledge-based protocol editor.

One extension to this formalisation process is to include representations of regulations and standards of authorities such as the US Food Drugs Agency or guidelines produced by the EEC[21] and relevant professional bodies[22]. Then we need to add another component to the knowledge base:

Standards description of essential components of a clinical trial protocol

The contribution these can make to the formal model can be summarised in a further variant of (1) and (2) above. We can include the standards in the knowledge base and use them deductively to produce symbolic representations of components of the protocol. This corresponds to (3) below

(3) Oncology $\cup$ Safety $\cup$ Standards $\cup$ TherapyPlan $\vdash$ ProtocolSafety $\cup$ ProtocolStructure

In order to implement such a model in a protocol editing system, it will be necessary first to undertake

a detailed analysis of the appropriate regulatory documents and guidelines.

There are obvious benefits in reusing knowledge bases for different activities and much related work in AI is currently being undertaken. To an individual clinician who needs to design protocols on an infrequent basis, software aids will provide significant time savings. If these aids include the ability to generate individual clauses concerning adverse events and other aspects of safety then compliance with the requirements of authorities such as the FDA, the guidelines of the EEC, and the recommendations of professional bodies are an additional benefit.

Acknowledgements
The authors gratefully acknowledge discussions with Professor Adrian Harris and Dr Nicola Stoner (ICRF Clinical Oncology Unit, Churchill Hospital, Oxford, UK) and suggestions for improvement from anonymous referees.

References

1 Lintzelman DK, Dittus RS, Miller ME, Tierney WM. Requiring physicians to respond to computerised reminders improves their compliance with preventive care reminders. J Gen Int Med. 1993; 8 (6): 311-317.

2 Kent DL, Shortliffe EH, Carlson RW, Bischoff MB, Jacobs CD. Improvements in data collection through physician use of a computer-based chemotherapy treatment consultant. J Clin Onc. 1985; 3 (10): 1409-1417.

3 Hammond P, Harris A L, Das SK, Wyatt JC. Safety and decision support in oncology. Methods of Information in Medicine. 1994; 33(4): 371-381.

4 Hammond P, Sergot MJ. Computer support for protocol-based treatment of cancer. Journal of Logic Programming. 1995; in press.

5 Wyatt J, Altman D, Heathfield H, Pantin C. Development of Design-a-Trial, a knowledge-based critiquing system for authors of clinical trial protocols. Comp Prog Meth Biomed. 1994; 43: 283-291.

6 Musen MA, Combs DM, Walton JD, Shortliffe EH and Fagan LM. OPAL: Toward the computer-aided design of oncology advice systems. Knowledge Systems Laboratory Memo KSL-86-49, Medical Computer Science, Stanford University, Stanford, California 94305-5479. 1986.

7 Fischer DS, Knobf MT, Durivage HJ. The cancer chemotherapy handbook. 4th Edition. Mosby, St Louis. 1993.

8 Protocol BO03. A randomised trial of two chemotherapy regimens in the treatment of operable osteosarcoma, MRC. 1986.

9 Alexander RB, Rosenberg SA: Tumour necrosis factor: Clinical applications. In: De Vita VT, Hellman S, Rosenberg SA (eds): Biological Therapy of Cancer. Philadelphia, JB Lippincott Co. 1991: 378-392.

10 Rowinsky EK, Gilbert MR, McGuire WP, et al: Sequences of taxol and cisplatin: A phase 1 trial and pharmacologic study. J Clin Oncol. 1991; 9:1692-1703.

11 Seymour MT, Slevin ML. A randomised trial of 5FU and leucovorin versus 5FU, leucovorin and interferon for advanced colorectal adenocarcinoma. Medical Research Council. UK.

12 Gallaire H, Minker J. (Eds.), Logic and Databases. New York: Plenum. 1978.

13 COSSH Control of substances hazardous to health Regulations. Approved Code of Practice. Health & Safety Commission, UK. 1988.

14 Mead GM, Kaye SB. A MRC/EORTC collaborative randomised trial of BEP/EP vs BOP/VIP-B with or without G-CSF in the management of poor prognosis metastatic and extragonadal teratoma. Medical Research Council. UK.

15 Langlotz C, Fagan L, Tu S, Williams J, Skic B. ONYX: an architecture for planning in uncertain environments. 1984.

16 Swartout WR. A Digitalis Therapy Adviser with explanations. MIT Laboratory for Computer Science TR-176, Cambridge, MA. 1977.

17 Swartout WR. A Digitalis Therapy Adviser with explanations. Proc of 5th IJCAI, Cambridge, MA. 1977.

18 Swartout WR. XPLAIN: a system for creating and explaining expert consulting programs. Artificial Intelligence. 1983; 21: 285-325.

19 Musen MA. Automated generation of model-based knowledge-acquisition tools. Pitman, London. 1989.

20 Van der Lei J, Musen MA. A model for critiquing based on automated medical records. Comput Biomed Res, 1991; 24: 344-378.

21 EC CPMP Working Party on Efficacy of Medicinal Products. Good Clinical Practice for Trials on Medicinal Products in the European Community. Pharmacology and Toxicology. 1990; 67: 361-372.

22 Cancer Research Campaign. Standard Operating Procedure for Protocol Preparation. Unpublished document. 1993.